Developments in Stochastic Coupled Cluster Theory: The initiator approximation and application to the Uniform Electron Gas

James S. Spencer
Department of Physics, Imperial College London, Exhibition Road, London, SW7 2AZ, United Kingdom and Department of Materials, Imperial College London, Exhibition Road, London, SW7 2AZ, United Kingdom

Alex J. W. Thom*
University Chemical Laboratory, Lensfield Road, Cambridge CB2 1EW, United Kingdom and Department of Chemistry, Imperial College London, Exhibition Road, London, SW7 2AZ, United Kingdom

(Dated: November 19, 2015)

We describe further details of the Stochastic Coupled Cluster method and a diagnostic of such calculations, the shoulder height, akin to the plateau found in Full Configuration Interaction Quantum Monte Carlo. We describe an initiator modification to Stochastic Coupled Cluster Theory and show that initiator calculations can be extrapolated to the unbiased limit. We apply this method to the 3D 14-electron uniform electron gas and present complete basis set limit values of the CCSD and previously unattainable CCSDT correlation energies for up to $r_s = 2$, showing a requirement to include triple excitations to accurately calculate energies at high densities.

arXiv:1511.05752v1 [physics.chem-ph] 18 Nov 2015

I. INTRODUCTION

In a recent letter[1], we described a novel way of formulating Coupled Cluster theory stochastically. The previous exposition was limited to calculations on a single processor, but was shown capable, within arbitrarily small error bars, of reproducing exact Coupled Cluster results for a range of molecular systems while requiring lesser computational resources than the full exact calculations.

Coupled Cluster theory, though born in the field of nuclear quantum mechanics, has become astoundingly successful in the field of quantum chemistry. After the pioneering work of Čížek and Paldus [2, 3] brought it to the quantum chemists' attention, and its use has grown such as to be the 'gold-standard' of quantum chemistry[4]. Its success stems, in essence, from the ability of CCSD(T) (Coupled Cluster Singles and Doubles with perturbative Triples)[5] to reproduce energetics and properties of a wide range of molecules to 'chemical accuracy', i.e. within 1 kJ mol^{-1}, and relative ease of formulation, allowing it to be implemented in a range of widely used computational chemistry software. It has not, however, become ubiquitous within electronic structure theory owing to a high scaling ($\mathcal{O}(N^7)$ for CCSD(T)[53]) with number of electrons N. Numerous local approximations[6, 7] have reduced this scaling, even to the level of linear with system size[8], though these have yet to become widespread, and often the onset of the reduced scaling is only in the regime of large systems. For greater accuracy one can turn to higher levels of truncation of the theory, with full triples, CCSDT, or full triples and quadruples, CCSDTQ, having scalings $\mathcal{O}(N^8)$ and $\mathcal{O}(N^{10})$ respectively. This higher accuracy is at the cost of their being vastly more complicated to implement, and their scaling makes them of little current use beyond benchmarking the energetics of very small molecules. It is in going beyond this regime that we believe the novel stochastic Coupled Cluster theory to be of great use owing to both its simplicity of implementation and parallelizability.

Within the paradigm provided by stochastic methods, there has been a great resurgence in their applications to quantum chemical methods. The highly acclaimed work of Booth, Alavi and co-workers[9, 10] in formulating Full Configuration Interaction Quantum Monte Carlo (FCIQMC) as a stochastic implementation of FCI was the impetus for genesis of the stochastic formulation of Coupled Cluster theory, and the common elements of the methods allow algorithmic developments to be easily transferred between them. In particular, we shall focus on the 'initiator approximation' of Cleland *et al.*[11–13] has allowed vastly larger systems to be studied than conventional diagonalization and both work on new properties (explicit correlation[14], density matrices[15], forces[16], and excited states[17, 18]) and new systems (from the Uniform Electron Gas[19, 20] and the Hubbard model[21], to real crystalline solids[22, 23]).

In this paper, we elucidate more algorithmic details of the CCMC method, as well the inclusion of the 'initiator approximation' based upon an analogous development within FCIQMC[11]. These advances allows the application of high levels of Coupled Cluster theory to problems much larger than previously possible with full Coupled Cluster theory. In Section II we describe the basic algorithm in detail and provide examples of the Monte Carlo steps involved and denote this algorithm Coupled Cluster Monte Carlo (CCMC). The population dynamics of CCMC are explored in Section III, followed by more details about the critical plateau height in Section IV in

*Electronic address: ajwt3@cam.ac.uk

which we propose "shoulder plots" as a measure of the critical difficulty of a system, and show how this varies in some molecular systems. Section V introduces the initiator approximation for CCMC and Section VI details how the systematic error inherent to the initiator approximation can be extrapolated to zero. The behaviour of the CCMC method is explored in these sections using the neon atom and nitrogen molecule, going to six-fold excitations (i.e. CCSDTQ56). The use of quantum chemistry methods in extended systems has been the subject of much recent interest[24–28] and as a further demonstration of the method, we calculate the CCSD and CCSDT energies of the uniform electron gas in Section VII. We find a substantial impact from triple excitations at realistic densities. Finally we offer some concluding remarks and future perspective in Section VIII.

II. COUPLED CLUSTER MONTE CARLO

Key to the success of Monte Carlo techniques is the ability to reduce computational effort by converting sums and integrals over extremely large spaces into a series of discrete samples which approximate the full calculation to arbitrary accuracy with increasing numbers of samples. In this section we cast the Coupled Cluster equations in such a form as can be sampled with Monte Carlo techniques, and show how by parameterization as discrete objects in excitation space the Coupled Cluster equations may be easily approximated and solved.

The space in which we shall represent single reference Coupled Cluster theory is that of *excitors* which act with respect to a *reference* Slater determinant. Given an orthonormal set of $2M$ spin-orbitals, we partition them into a set of N which are occupied in the reference determinant, denoted $\phi_i, \phi_j, \dots \phi_n$, and $2M - N$ unoccupied, or virtual orbitals, denoted $\phi_a, \phi_b, \dots \phi_f$. The complete space of N-electron Slater determinants in this basis has size $\binom{2M}{N}$ and will be denoted by $D_{\mathbf{m}}$ where $\mathbf{m}$ is an N-vector listing the orbitals occupied in a given determinant. Given the occupied/virtual partitioning, we may also represent all possible determinants with respect the the reference (which we will denote D_0) by listing the occupied orbitals removed and virtual orbitals added to the determinant, such that D_{ij}^{ab} represents the determinant where i and j in the reference have been replaced by a and b. Further we shall denote by $\hat{a}_{ij}^{ab}$ the *excitation operator* or *excitor* which performs this process, $D_{ij}^{ab} = \hat{a}_{ij}^{ab} D_0$. The excitors behave such that it is not possible to excite from or to an orbital multiple times, e.g. $\hat{a}_i^a \hat{a}_i^b = 0$ and $\hat{a}_i^a \hat{a}_j^a = 0$. By applying each of the $\binom{2M}{N}$ possible excitors to the reference determinant, the whole of determinant space may be generated, and there is a one-to-one correspondence between excitors and determinants; we may therefore denote the excitors by the determinant they would create, $D_{\mathbf{i}} = \hat{a}_{\mathbf{i}} D_0$.

These excitors can be used to parameterize all possible N-particle wavefunctions in this basis in a number of ways. A simple example would be $\Psi_{\mathrm{FCI}} = \sum_{\mathbf{i}} c_{\mathbf{i}} \hat{a}_{\mathbf{i}} D_0$, which would correspond a Full Configuration interaction type wavefunction where each determinant has its own coefficient in the expansion. Coupled Cluster theory uses instead an exponential *Ansatz* for the wavefunction, $\Psi_{\mathrm{CC}} = e^{\hat{T}} D_0$, where we define $\hat{T} = \sum_{\mathbf{i}} t_{\mathbf{i}} \hat{a}_{\mathbf{i}}$. This seemingly complicated parameterization is used owing to its desirable property of remaining size-consistent even if the sum of excitors is restricted to a limited level of excitation. In essence this is due to the fact that despite a truncation, the exponential ensures that the wavefunction can contain contributions from determinants at all excitation levels.

To determine the parameters $\{t_{\mathbf{i}}\}$, the projected Schrödinger equation is solved; i.e. for all $\mathbf{m}$, we solve

$$\langle D_{\mathbf{m}} | \hat{H} - E | \Psi_{\mathrm{CC}} \rangle = 0. \tag{1}$$

In general this may be expressed in an iterative form, beginning with a guess for all $\{t_{\mathbf{i}}\}$ and E, and iterating until convergence. The complexity in this arises from the expansion of the exponential

$$e^{\hat{T}} D_0 = \left[1 + \sum_{\mathbf{i}} t_{\mathbf{i}} \hat{a}_{\mathbf{i}} + \frac{1}{2} \sum_{\mathbf{ij}} t_{\mathbf{i}} t_{\mathbf{j}} \hat{a}_{\mathbf{i}} \hat{a}_{\mathbf{j}} + \frac{1}{3!} \sum_{\mathbf{ijk}} t_{\mathbf{i}} t_{\mathbf{j}} t_{\mathbf{k}} \hat{a}_{\mathbf{i}} \hat{a}_{\mathbf{j}} \hat{a}_{\mathbf{k}} + \dots \right] D_0. \tag{2}$$

Instead of explicitly rearranging the equations (1) (or the equivalent more convenient equations, $\langle D_{\mathbf{m}} | e^{-\hat{T}} (\hat{H} - E) | \Psi_{\mathrm{CC}} \rangle = 0$) in an iterative form, which is the means by which many conventional implementations work), we note that solutions to the Coupled Cluster equations must also satisfy

$$\langle D_{\mathbf{m}} | 1 - \delta\tau (\hat{H} - E) | \Psi_{\mathrm{CC}} \rangle = \langle D_{\mathbf{m}} | \Psi_{\mathrm{CC}} \rangle, \tag{3}$$

where $\delta\tau$ is some small positive number[54], giving a form reminiscent of projector methods. This is now almost in the form of an iterative method:

$$\langle D_{\mathbf{m}} | \Psi_{\mathrm{CC}} \rangle - \delta\tau \langle D_{\mathbf{m}} | \hat{H} - E | \Psi_{\mathrm{CC}} \rangle = \langle D_{\mathbf{m}} | \Psi_{\mathrm{CC}} \rangle, \tag{4}$$

except that the right hand side has components from many different $t_{\mathbf{i}}$ amplitudes, so it is not clear how to perform the iteration. Noting that the projection of Ψ_{CC}

on to a single determinant always contains the amplitude of the excitor for that determinant plus terms involving multiple amplitudes,

$$\langle D_{\mathbf{m}}|e^{\hat{T}}D_0\rangle = t_{\mathbf{m}} + \mathcal{O}(\hat{T}^2), \tag{5}$$

we may cancel out the higher order terms identically on the left and right hand sides of (4), leaving

$$t_{\mathbf{m}} - \delta\tau\langle D_{\mathbf{m}}|\hat{H} - E|\Psi_{\mathrm{CC}}\rangle = t_{\mathbf{m}}. \tag{6}$$

We may write this as an iteration from time τ to $\tau + \delta\tau$,

$$t_{\mathbf{m}}(\tau) - \delta\tau\langle D_{\mathbf{m}}|\hat{H} - E|\Psi_{\mathrm{CC}}(\tau)\rangle = t_{\mathbf{m}}(\tau + \delta\tau). \tag{7}$$

The solutions to the Coupled Cluster equations will also be solutions of this iteration procedure[55], though the evaluation of the second term in (7) is not trivial. Indeed, it turns out to be easier to sample this term. The most straightforward way of sampling these equations involves storing all $t_{\mathbf{i}}$ amplitudes as real numbers, and performing two stochastic processes.

The first is sampling Ψ_{CC}, by selecting randomly from all possible clusters in (2), e.g. $t_{\mathbf{i}}t_{\mathbf{j}}\hat{a}_{\mathbf{i}}\hat{a}_{\mathbf{j}}D_0$. Once selected the cluster is collapsed to form determinant, $t_{\mathbf{i}}t_{\mathbf{j}}\hat{a}_{\mathbf{i}}\hat{a}_{\mathbf{j}}D_0 = t_{\mathbf{i}}t_{\mathbf{j}}\hat{a}_{\mathbf{n}}D_0 = t_{\mathbf{i}}t_{\mathbf{j}}D_{\mathbf{n}}$. This process may involve some sign changes or even result in zero (if **i** and **j** excite from the same occupied or to the same virtual orbitals).

The second stochastic process is the sampling the action of the Hamiltonian. For each $D_{\mathbf{n}}$ generated from sampling Ψ_{CC}, we could enumerate all possible $D_{\mathbf{m}}$, evaluating $\langle D_{\mathbf{m}}|\hat{H} - E|D_{\mathbf{n}}\rangle$ and updating $t_{\mathbf{m}}$ accordingly. As the Hamiltonian only connects up to single and double excitations from a determinant, we may instead sample this process by randomly picking $D_{\mathbf{m}}$ as a single or double excitation of $D_{\mathbf{n}}$, and updating $t_{\mathbf{m}}$ appropriately. While we have found such an algorithm which stores all $t_{\mathbf{n}}$ to well reproduce small Coupled Cluster calculations, it is not efficient, as many clusters selected have extremely small amplitudes, requiring large numbers of samples.

A more efficient route is to *discretize* the amplitudes as has been done in the FCIQMC method[9]. In analogy to Anderson[29], we shall represent amplitudes of excitors by sets of *excitor particles* or *excips*[56]. Each excip represents a unit weight and is given a positive or negative sign, and located at an excitor, $\hat{a}_{\mathbf{n}}$. As the simulation proceeds, excips are created or destroyed at excitors according to some simple rules, and the mean (averaged over a number of iterations) signed number of excips at a given excitor will be taken to represent its (unnormalized) amplitude.

To sample the action of (7), we make very similar steps to the algorithm using real numbers just described. Firstly a cluster out of Ψ_{CC} is selected randomly, and collapsed into a determinant $D_{\mathbf{n}}$: since $\hat{T}$ is represented by populations of excips at different excitors, a cluster is formed by randomly selecting a number of excips (which could be zero). Clusters in Ψ_{CC} formed from excitors with no excips at them would have no amplitude, and so may be safely ignored. The cluster is collapsed by taking the product of the excitors at which each of the chosen excips is located to give determinant $D_{\mathbf{n}}$. The amplitude of the cluster is given by the products of populations of excips at each of the excitors in the cluster. As the spawning probability must be proportional to this, we have chosen to select each cluster with a probability proportional to its amplitude, and make one spawning attempt per cluster. Then, much as in FCIQMC, we follow the following steps:

1. *Spawning.* From each $D_{\mathbf{n}}$ generated we pick a random connected single or double excitation, $D_{\mathbf{m}}$, and create a new excip there with probability proportional to $|\delta\tau H_{\mathbf{mn}}|$ and appropriate sign.

2. *Birth/Death.* From each $D_{\mathbf{n}}$ generated, with probability proportional to $|\delta\tau(H_{\mathbf{nn}} - E)|$ we create an opposite signed excip at $\hat{a}_{\mathbf{n}}$. Here E has yet to be defined, but may be taken as an arbitrary constant. Its later role will be to control population growth. Because $D_{\mathbf{n}}$ may have been formed from a product of excitors, it is probable that there are no excips at the excitor $\hat{a}_{\mathbf{n}}$, and so, in contrast to FCIQMC, we must create an oppositely signed excip, rather than simply killing an existing one. If the cluster consisted of a single excip, the opposite signed excip will cancel this out.

3. *Annihilation.* Finally, after these two processes have been done for as many clusters as required to be sampled, the new list of excips is sorted and any opposite signed pairs of excips on the same excitor removed.

In deterministic Coupled Cluster theory, E is a parameter which iteratively converges to the Coupled Cluster energy. Here we shall replace it with a parameter S which will shift the diagonal elements of the Hamiltonian and may be updated periodically after a number of cycles to allow the total population of excips to be controlled[9, 30].

A. Normalization

The careful reader may have noticed that above formalism introduces an imbalance in the description of the reference and the excitor space. Indeed, the prescription of intermediate normalization, where the overlap of the reference with the Coupled Cluster wavefunction is fixed at unity, makes it rather difficult to satisfy (4) for $D_{\mathbf{m}} = D_0$ in terms of an iteration. Instead, it is convenient to introduce a further variable, N_0, to act as the normalization constant,

$$|\Psi_{\mathrm{CC}}\rangle = N_0 e^{\frac{\hat{T}}{N_0}}|D_0\rangle. \tag{8}$$

The inclusion of N_0 within the exponential allows the discretization of $\hat{T}$ to produce effectively fractional populations on excitors which is essential for the exponential

to converge. With this normalization, the number of excips at the reference is given by N_0, and can thus be varied in the iteration satisfying (4). In this manuscript we will refer to these particles on the reference as excips, though they are not strictly the same, and cannot be used as a constituent excip of a cluster; as such they are (all) always in a cluster as they normalize it multiplicatively.

B. Energy estimators

There are two independent estimators of the Coupled Cluster energy available in this formalism. Firstly, the time-averaged value of the shift is, at convergence, an estimator for the energy. Secondly the energy may be estimated by projection as

$$E_{\text{proj}} = \frac{\langle D_0|\hat{H}|\Psi_{\text{CC}}\rangle}{\langle D_0|\Psi_{\text{CC}}\rangle} \tag{9}$$

This expression could be computed exactly, as the contributions from only singly and doubly excited determinants are relevant, but would require an iteration through all possible pairs of single excitors to evaluate their contribution to the doubly excited determinants. Instead, we sample both of these contributions while sampling the cluster wavefunction.

C. Sampling

Sampling forms the core of this method, and the implementational details of this are crucial to an efficient calculation. Here we describe the different sampling schemes used in our implementation. Though we have not performed exhaustive tests on different sampling methods, these have appeared the most effective in our calculations.

1. *Sampling the wavefunction.* At each timestep we must sample the exponential expansion of the wavefunction, and we do this by randomly selecting a cluster — a single term consisting of a product of specific amplitudes and excitors — from the expansion of the exponential. We have found that setting the number of samples taken to be the same as the number of excips leads to stable simulations. If the number of samples does not scale at least linearly with the number of excips, the selection probabilities of the clusters become increasingly small and lead to population blooms which destabilize the simulation.

   We select the size of cluster, s, to be generated so that the probability of selecting a cluster with s excitors decays exponentially with s. Each cluster is generated by selecting s excips randomly from the list of excips, and we have found it best to bias the selection so that the probability of selecting an excip is proportional to the total number of excips at that excitor.

2. *Sampling the spawning.* In principle the spawning step samples the action of the Hamiltonian upon the collapsed cluster. A full sampling would enumerate all connected determinants and attempt to spawn on each one-by-one. We have followed the route of Booth *et al.*[9] in choosing a minimalist sampling of a single attempted spawn from each cluster. This comes with the drawback that as the number of connected determinants increases (say with system size or basis), the timestep must be correspondingly reduced to keep this step stable. Some brief tests of increasing the number of spawnings from each cluster indicate this has the same effect as dividing the timestep by this number (reducing the effective timestep).

3. *Sampling the projected energy.* The explicit evaluation of the numerator of (9) is complicated by pairs of single-excitations. Instead, each time a cluster is generated, its contribution to the numerator is accumulated (appropriately unbiased by its generation probability).

D. A Representative Example

Let us consider a spin-free system with occupied orbitals i, j, k and virtual orbitals a, b, c. In addition to the reference there are 19 possible excitors, which we shall denote with subscripts and superscripts. For example, $_{ik}^{ab}$, denotes the excitor which excites electrons from i and k to orbitals a and b.

Let us suppose we are performing Coupled Cluster theory truncated at singles and doubles (a rather futile truncation since there is only one triple excitation), i.e. CCSD. At time τ in the simulation, we have reached the situation at the top of Fig. 1(a), with 20 excips, 10 being at the reference. We shall accumulate various generation probabilities, p_{X}, through this example which will be used to unbias our sampling. We sample with the same number of samples as excips (hence $p_{\text{sel}} = 20$ for this timestep). We shall also adopt a scheme where the likelihood of forming a cluster is exponentially decaying with probability $\frac{1}{2}$. Since the largest cluster with any effect is of size three, we stop at that point and add in the remaining probabilities of selection for larger clusters to that level. Let us also suppose that we have not yet modified the shift from its default $S = E_{\text{HF}}$. The following four examples of samples and their evolution are illustrated in Fig. 1(a)–1(d).

(a) We decide with probability $p_{\text{size}} = \frac{1}{2}$ to select a cluster with zero excips. There is only one choice of zero-cluster, so we set $p_{\text{clust}} = 1$. The action of this on the reference is merely $\hat{1}|D_0\rangle = |D_0\rangle$, and it has total amplitude $A = N_0$. There

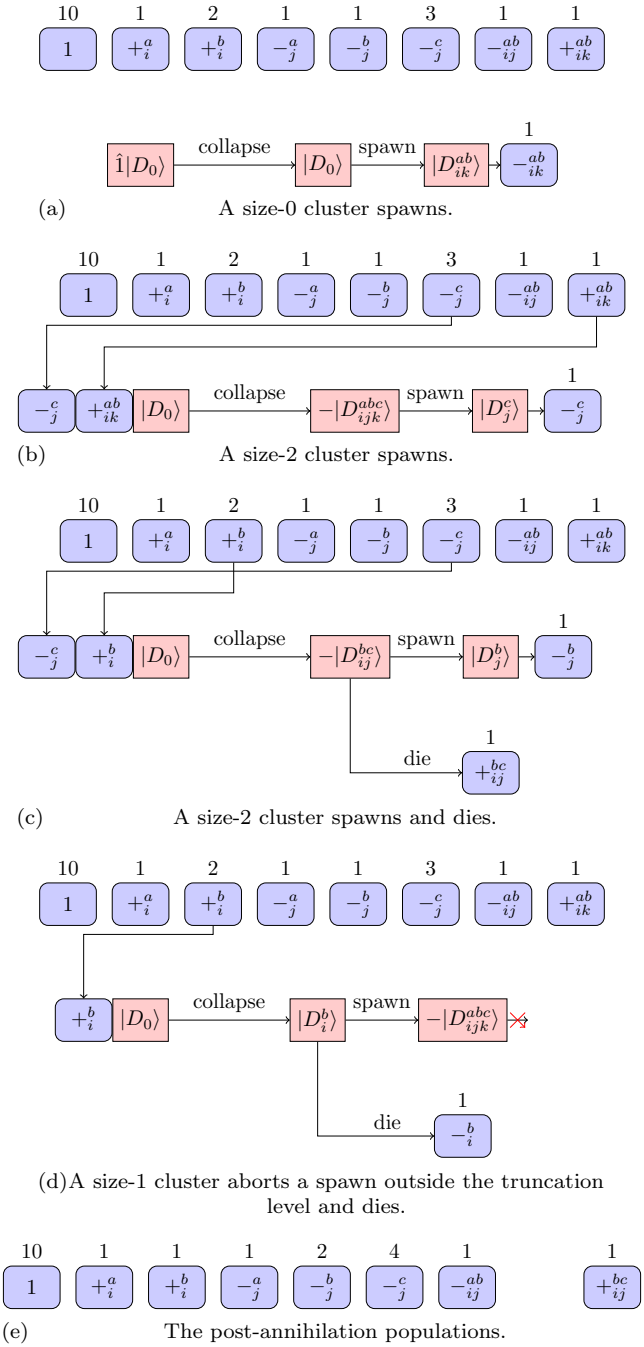
10 1 2 1 1 3 1 1
collapse
spawn
die
(a) A size-0 cluster spawns.
(b) A size-2 cluster spawns.
(c) A size-2 cluster spawns and dies.
(d) A size-1 cluster aborts a spawn outside the truncation level and dies.
(e) The post-annihilation populations.

FIG. 1: A representative example of the cluster selection, spawning and death steps of the CCMC algorithm as explained in the main text. Rounded blue boxes represent excitors, and rectangular blue boxes represent determinants.

are eighteen possible single and double excitations available, and we pick D_{ik}^{ab} with $p_{\text{excit}} = \frac{1}{18}$. $\hat{a}_{ik}^{ab}|D_0\rangle = |D_{ik}^{ab}\rangle$, so no sign change will be needed. We create a new excip at $^{ab}_{ik}$ with probability $\delta\tau|\langle D_{ik}^{ab}|\hat{H}|D_0\rangle A|/(p_{\text{sel}}p_{\text{size}}p_{\text{clust}}p_{\text{excit}})$, and sign $-\text{sgn}(\langle D_i^b|\hat{H}|D_0\rangle)$. Let's say we create a negative excip at $^{ab}_{ik}$. Death occurs with probability $|(\langle D_0|\hat{H}|D_0\rangle - S)\delta\tau A|/(p_{\text{sel}}p_{\text{size}}p_{\text{clust}})$. Here as $E_{\text{HF}} = S$, there is no death.

(b) We decide with probability $p_{\text{size}} = \frac{1}{8}$ to pick a cluster of size 2. $-_j^c$ is selected with probability $p_1 = \frac{3}{10}$ and $+_{ik}^{ab}$ is selected with probability $p_2 = \frac{1}{10}$, so $p_{\text{clust}} = (2!)p_1p_2$, the 2! accounting for the number of ways in which the cluster could have been chosen. The amplitude of the cluster is given by $A = N_0\frac{-3}{N_0}\frac{1}{N_0}$. The cluster $\hat{a}_j^c\hat{a}_{ik}^{ab}$ collapses to $-|D_{ijk}^{abc}\rangle$ when applied to the reference. We first spawn, picking $|D_j^c\rangle$ with probability $p_{\text{excit}} = \frac{1}{18}$.

$\hat{a}_j^c|D_0\rangle = -|D_j^c\rangle$ so we will pick up an extra negative if we spawn there, which we do with probability $\delta\tau|\langle D_j^c|\hat{H}|D_{ijk}^{abc}\rangle A|/(p_{\text{size}}p_{\text{clust}}p_{\text{excit}})$. The signs take a little more care. Let's say that $\langle D_j^b|\hat{H}|D_{ijk}^{abc}\rangle$ is negative, giving a collective five negatives, resulting at a negative excip at ^c_j. Death would occur with probability $|(\langle D_{ijk}^{abc}|\hat{H}|D_{ijk}^{abc}\rangle - S)\delta\tau A|/(p_{\text{sel}}p_{\text{size}}p_{\text{clust}})$, except that in CCSD, we do not store excips for triples, so it is ignored.

(c) We decide with probability $p_{\text{size}} = \frac{1}{8}$ to pick a cluster of size 2. $-_j^c$ is selected with probability $p_1 = \frac{3}{10}$ and $+_i^b$ is selected with probability $p_2 = \frac{2}{10}$, giving $p_{\text{clust}} = (2!)p_1p_2$. The amplitude is $A = N_0\frac{-3}{N_0}\frac{2}{N_0}$. The cluster $\hat{a}_j^c\hat{a}_i^b$ collapses to $|D_{ij}^{bc}\rangle$ when applied to the reference. We pick $|D_j^b\rangle$ with probability $p_{\text{excit}} = \frac{1}{18}$. $\hat{a}_j^b|D_0\rangle = -|D_j^b\rangle$ so we will pick up an extra negative for spawning, which has probability $\delta\tau|\langle D_j^b|\hat{H}|D_{ij}^{bc}\rangle A|/(p_{\text{size}}p_{\text{clust}}p_{\text{excit}})$. A positive value of $\langle D_j^b|\hat{H}|D_{ij}^{bc}\rangle$ results in a negative excip being spawned at ^b_j. Death occurs with probability $|(\langle D_{ij}^{bc}|\hat{H}|D_{ij}^{bc}\rangle - S)\delta\tau A|/(p_{\text{sel}}p_{\text{size}}p_{\text{clust}})$, creating a positive excip at $^{bc}_{ij}$ because A is negative.

(d) We decide with probability $p_{\text{size}} = \frac{1}{4}$ to pick a cluster of size 1. $+_i^b$ is selected with probability $p_1 = \frac{2}{10}$, giving $p_{\text{clust}} = (1!)p_1$. The amplitude is $A = N_0\frac{2}{N_0}$. The cluster $\hat{a}_i^b$ collapses to $|D_i^b\rangle$ when applied to the reference. We pick spawning excitation $|D_{ijk}^{abc}\rangle$ with probability $p_{\text{excit}} = \frac{1}{18}$. This is a triple excitation and outside our truncation so it is discarded. Death occurs with probability $|(\langle D_i^b|\hat{H}|D_i^b\rangle - S)\delta\tau A|/(p_{\text{sel}}p_{\text{size}}p_{\text{clust}})$, creating a negative excip at ^b_i.

Let us assume that all the other samplings result in no death or spawning. The resultant list of new excips is merged with the main list, annihilation occurs, and a new step begins (Fig. 1(e)). We note that in this implementation much of the amplitude information cancels between A and p_{clust}, resulting an a relatively efficient sampling.

III. POPULATION DYNAMICS

The behaviour of the such algorithms has been previously explored in the context of FCIQMC[31], and follows the following phases as shown in Fig 2:

1. *Initial Growth.* We beginning with a population of excips at the reference. By initially setting $S = E_{\text{HF}}$ we ensure that the excips at the reference do not die. Gradually, the excitor space is populated with excips spreading outwards from the reference. As time progresses and the populations grow, annihilation events become more common.

2. *Annihilation Plateau.* At a system-dependent number of excips, annihilation events become equal to spawning events, and the population ceases to grow, but remains approximately constant. The relative signs of the populations at excitors where there are many excips changes over this time until it resembles the sign structure appropriate to the ground state wavefunction (as parameterized by excitation amplitudes).

3. *Exponential Growth.* The sign structure of the excips is such that spawning once again dominates over annihilation, and the total number of excips grows exponentially with growth rate proportional to $E - E_{\text{CC}}$, where E_{CC} is the Coupled Cluster energy.

4. *Population Control.* S is now dynamically modified to control the population and limit growth, with more negative values of S reducing population growth.

Before the Annihilation Plateau, the total rate of growth is faster than that at the reference, and if one attempts to control the population growth with the shift, the population at the reference decreases to zero, causing the simulation to fail. Below the plateau, in common with FCIQMC, we also expect the sign structure of the solution to be wrong[31], and so the plateau height becomes a crucial feature in the usability of Stochastic Coupled Cluster.

We note that, as we begin simulations with a relatively large population of excips at the reference, often the rate of annihilation does not exactly match that of growth and the population growth moves through a shoulder rather than a plateau (see figure 2). This is caused by effective overshooting the plateau phase by starting with more excips. To determine whether the system has reached the region of exponential growth, we note that in this region, the population of at all excitors should grow exponentially. A suitable alternative is therefore to check whether the population at the reference is growing at the same rate as that of the total population. One this has been satisfied, the shift, S, may be allowed to vary to control the population.

IV. THE PLATEAU

The stability of a CCMC calculation requires reaching a plateau or shoulder in the particle growth. Discerning such a feature by eye is not an easy task, so we sought a more rigorous definition. Before the plateau, a shift based on the (faster) total particle growth rate would cause a decrease in the population at the reference, and likely lead the calculation to fail as this population tends to zero. The post-plateau stability arises from the fact that the growth rate of reference particles is at least as fast as those of the total particles in the system, and so

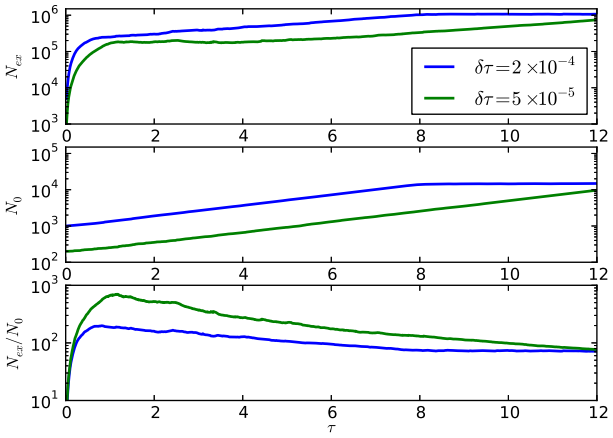

(a)

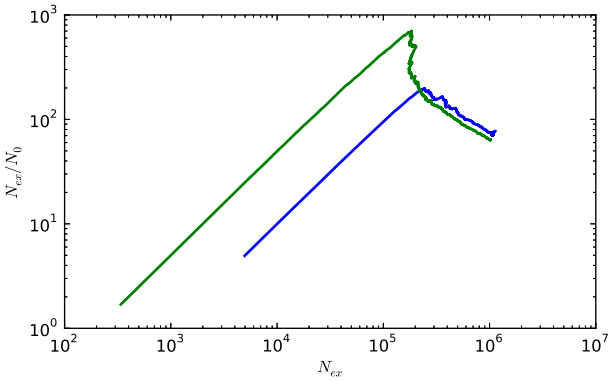

(b)

FIG. 2: Ne cc-pVQZ CCSDTQ calculations starting with different initial particle numbers at the reference and different timesteps. (a): With a carefully chosen low timestep and initial population, a plateau is visible. An increased timestep and initial population overshoots the plateau but has a shoulder. The lower panel shows a maximum of the particle ratio at the position of the shoulder and plateau. (b): 'Shoulder plots' allow shoulder height to be read off easily and calculations compared.

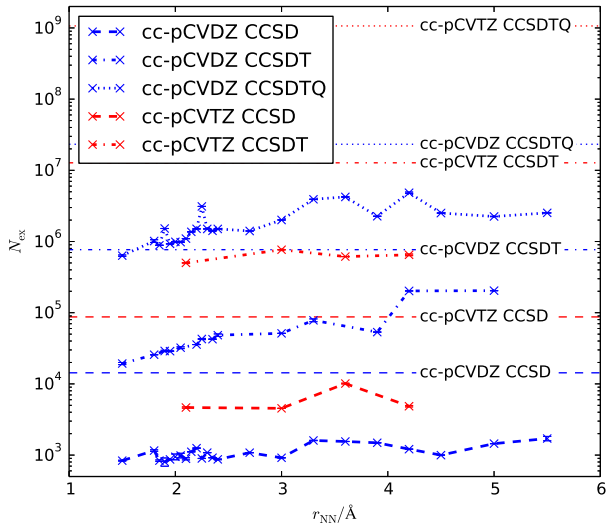

FIG. 3: Shoulder heights for the N_2 molecule at different bond lengths, basis sets and truncations. All electrons were correlated, and L_z symmetry enforced. Hilbert space sizes are shown by the horizontal lines.

turning on a shift will slow these rates equally and the populations will stabilize. On this basis, we suggest that the ratio of total particles to reference particles should be a useful measure of these rates of growth, and we shall call this the *particle ratio*. In figure 2 we plot a typical CCMC calculation showing a plateau and a second calculation on the same system with different timestep and larger initial population. The calculation with larger initial population overshoots the plateau, but still reaches a stable growth phase. The particle ratio shows clearly the location of the plateau by moving from increasing to downward; the overshooting calculation shows the same behaviour corresponding the shoulder in its growth.

From some experience with comparing such plots, it proves convenient to plot the particle ratio not against the progress in imaginary time, but against the total number of particles, both axes being logarithmic. In such a plot the plateau appears as a sharp downward kink in an otherwise apparently linear progress. The usefulness of such 'shoulder' plots in comparing different calculations is explained in figure 2(b). There is still some potential uncertainty as to the definition of the shoulder height, but from experience it has proven to be useful to define this to be at the maximum of the shoulder plot, and this should be taken as an upper limit of the plateau height of the system. Such a definition is still subject to stochastic noise, and so in this paper we have quoted values which are the mean of the total excip populations with the ten largest particle ratios, and used the standard deviation of these as a measure of uncertainty. Calculations with total particle numbers before the shoulder are often unstable, and calculations after the shoulder are usually stable, though the stochastic nature of such calculations removes total certainty about this. A stable population (due to the control from the variable energy shift) appears in the shoulder plot as a tight blob at the end of the line due to (comparatively) small stochastic fluctuations in both the population on the reference and the total population.

A histogram approach has been previously used to automate detection of the plateau in FCIQMC[21]; the shoulder plot is equally amenable to automation and both approaches give similar values for simulations with clear plateaus. Shoulder plots are more reliable when the plateau is shorter in duration or tends towards a shoulder in the population growth.

As shown in the original CCMC paper[1], the heights

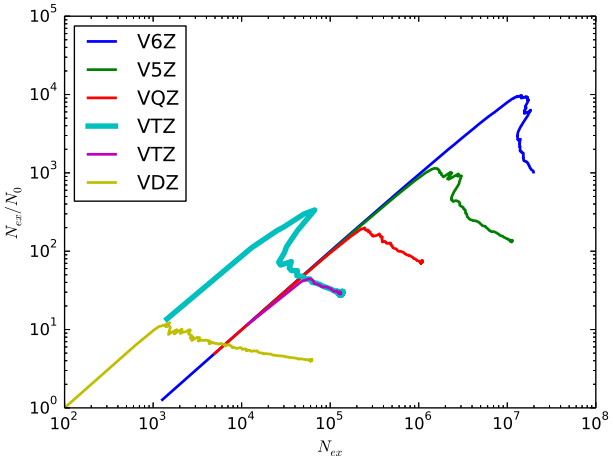

FIG. 4: All-electron CCSDTQ calculations on the neon atom using Dunning's cc-pVXZ basis sets[32, 33]. The shoulder is clearly discernible in each line, and its height increases with basis set cardinality. A vertical descent from the maximum is given by a clearly defined plateau, and a chevron-like peak is characteristic of a shoulder. As seen for cc-pVTZ, calculations started with more excips lead to a shoulder which can overestimate the position of a plateau. Calculations were started with 10^2 or 10^3 excips at the reference (where each line would cross $N_{\text{ex}}/N_0 = 10^0$).

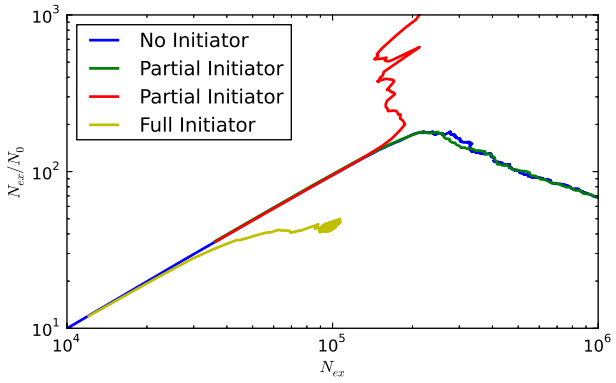

FIG. 5: Ne cc-pVQZ CCSDTQ calculations without initiator, and with partial and full initiator methods. The full initiator (yellow) method has a significantly lower shoulder height and is stable at even 100,000 excips whereas the partial (red) is unstable when the shift is engaged at 240,000 excips (before the plateau), and has a similar profile (green) to the non-initiator (blue) method.

of plateaux vary with both the size of Hilbert space and differing levels of correlation within a molecule; system-specific behaviour is also seen for FCIQMC calculations[9, 21]. Figure 3 shows the variation of plateau height with basis, excitation and bond length for the N_2 molecule, showing broadly that the fraction of the Hilbert space required is of a similar magnitude for different excitation levels and basis sets, and increases as the N−N bond is broken, though staying below 27%.

To see the effects of different basis sets, the neon atom, with only 10 electrons, is more amenable to study, and Figure 4 shows examples of shoulder plots for increasingly large bases. Plateau heights and correlation energies are given in Table I and decrease with truncation level, heading to less than 1% of the complete Hilbert space size for larger bases with quadruples truncation. Whilst for CCSD, therefore, the saving is relatively small, it is for higher excitation levels that the Monte Carlo sampling considerably reduces the storage required.

V. THE INITIATOR METHOD

With the insight that the plateau is exited only when the particles achieve a sign-coherent structure, Cleland et *al.* introduced what they call the 'initiator method'[11] to the FCIQMC algorithm (denoted i-FCIQMC). In its mature version it avoids the effects of sign-incoherence by restricting spawning. If a determinant has no walkers present, then the algorithm only allows a walker to be spawned there if there are more than a critical number, n_{add}, of walkers present at the determinant from which it was spawned. This modification of the algorithm dramatically alters the dynamics, removing the plateau phase and greatly stabilizing the calculations. This is, however, at the price of introducing a bias into the algorithm, dependent on the number of walkers. This is manifested as an error in the energy which systematically decreases with the number of walkers (becoming zero in the large walker limit).

Owing to the great similarity between Stochastic Coupled Cluster theory and FCIQMC, it is expected that the initiator method will confer similar benefits in Stochastic Coupled Cluster theory. As spawning occurs from clusters of excips rather than single walkers, the initiator method requires some modification. The obvious choices are for a cluster to be allowed to initiate (i.e. spawn onto an unoccupied excitor) if either one or all of its component excips come from an excitor whose population is $> n_{\text{add}}$. We denote these the partial and full initiator methods respectively. A representative sample of the behaviour of each of these schemes is shown in Figure 5. From this, and with the ideas of maintaining sign-coherence, we have settled upon requiring all excips within a cluster to come from excitors with populations $> n_{\text{add}}$ before allowing it to initiate (the full initiator method).

As the initiator calculation reduces the amount of spawning, it can be thought of as generally reducing the particle noise within a calculation. To investigate the effect of this, we note that in the regime of large numbers of particles, the variance of the numerator and denominator of the projected energy should decrease with $1/N_{\text{ex}}$. More formally, denoting the numerator of the projected

truncation	SD		SDT		SDTQ		SDTQ5		SDTQ56	
basis	H.S.	Shoulder	H.S.	Shoulder	H.S.	Shoulder	H.S.	Shoulder	H.S.	Shoulder
cc-pVDZ	400	33(1)	4680	194(2)	30654	594(4)	113550	1610(20)	259460	2660(10)
cc-pVTZ	1940	160(30)	60710	2900(300)	1.032×10^6	$2.27(1)\times10^4$	1.021×10^7	$1.90(5)\times10^4$	6.143×10^7	$9.6(2)\times10^5$
cc-pVQZ	6710	290(10)	4.128×10^5	7100(200)	1.412×10^7	$2.39(6)\times10^5$	2.81×10^8	$3.7(1)\times10^6$		
cc-pV5Z	17100	2500(400)	1.815×10^6	$7.9(2)\times10^4$	1.068×10^8	$1.57(6)\times10^6$	3.66×10^9			
cc-pV6Z	82400	$2.2(3)\times10^4$	1.636×10^7	$2.3(3)\times10^5$	1.731×10^9	$1.45(3)\times10^7$				

TABLE I: All-electron Ne Hilbert space (H.S.) sizes and shoulder heights. Hilbert space sizes were determined by Monte Carlo and only significant figures are shown. L_z symmetry was enforced except in the case of cc-pV6Z. Shoulder error estimates in the last significant digit are shown in parentheses.

energy at a given timestep to be $E_{\text{proj,numer}}(\tau)$, to enable comparison of different calculations, we may calculate

$$s^2 = \frac{\text{Var}[E_{\text{proj,numer}}(\tau)]_\tau \langle N_{\text{ex}}(\tau)\rangle_\tau}{\langle E_{\text{proj,numer}}(\tau)\rangle_\tau^2}, \qquad (10)$$

which we shall call the normalized relative variance of the projected energy. This is a substantial component of the eventual error estimate for a given calculation (though the actual errors are calculated through a blocking analysis[34]). A calculation with double the value of s^2 will have to be run for approximately twice as long to achieve the same estimate of error, and so it should be as low as possible to reduce the computational resources required in a calculation. For N_2 in a cc-pCVDZ basis, we show s^2 for both standard and initiator calculations in Figure 6. Around equilibrium, it can be seen that the variance is significantly lower for initiator calculations, and so using the initiator method should dramatically reduce the number of steps required to achieve a given error bar.

In Figure 7, we show data from a typical set of initiator calculations using CCSDT on the neon atom in a cc-pV6Z basis. As the population increases, the correlation energy found gradually converges to the true correlation energy. In this case the population required to for the initiator approximation's bias to have decayed is of the order of 5×10^5 psips, compared to a shoulder height of $2.3(3)\times10^5$ psips, so providing no benefit in reducing the number of psips needed. In all initiator CCMC calculations we have tried on systems in this paper (on the neon atom, N_2 molecule, and UEG), this pattern is repeated, and we see that a population of the order of the shoulder height is required to converge the initiator error. This is in contrast to in FCIQMC where the initiator adaptation substantially helps the sign problem allow much smaller populations to be used, and unfortunately we do not see any similar reductions in CCMC. The reduction in variance is, however, still beneficial.

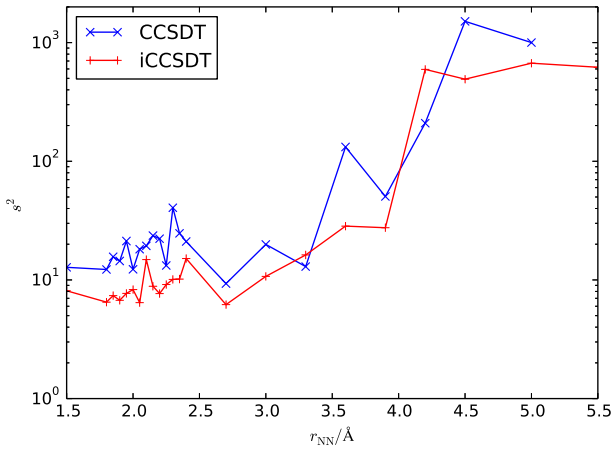

FIG. 6: The normalized relative variance of the projected energy, s^2, (see text for definition) for the N_2 molecule in a cc-pCVDZ basis. Calculations were performed with about 70,000 excips (for $r_{\text{NN}} \leq 2.4$ Å) and 200,000 otherwise. The initiator calculations show a generally reduced variance around equilibrium, leading to fewer steps being required to converge to a given error estimate.

VI. EXTRAPOLATION

In previous studies of the initiator approximation for FCIQMC, the initiator error has been seen to decay as total particle number increases, but there has been little known about the form of this decay. In appendix A, we show that the initiator error is expected to fall off as the inverse of the number of particles in iFCIQMC, and we shall use a similar form to extrapolate results for iCCMC, writing the energy as a function of number of excips, N_{ex}:

$$E_i(N_{ex}) = pN_{ex}^q + E_i(\infty),$$

where we will allow p, q, and $E_i(\infty)$ to be variable parameters. To test the form, we have performed CCSDT calculations on neon in a cc-pV6Z basis set, where the shoulder lies at $2.3(3)\times10^5$ excips out of a total space of 1.64×10^7 excitors, and we have been able to perform the exact CCSDT calculation with MOLPRO[35, 36] and MRCC[37]. For this system, iCCMC calculations were performed for increasing total excip populations, and projected energy estimators and errors determined using an automated blocking script[57]. As noted in appendix B, biases are introduced in calculations in which the shift varies greatly, and we have kept only calculations in which the standard deviation of the shift is lower

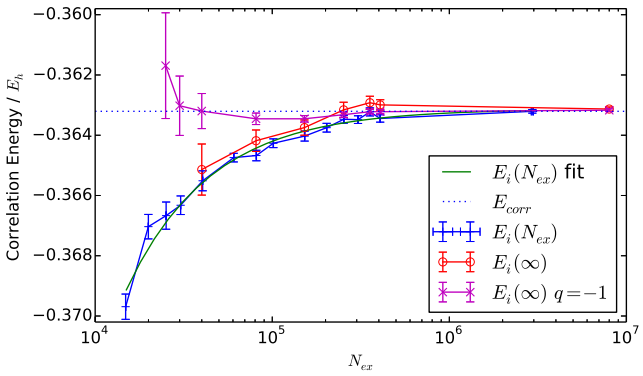

FIG. 7: CCSDT initiator energies and extrapolations for Ne cc-pV6Z. Plotted are the raw iCCMC energies (blue), $E_i(N_{ex})$, and various fits. The curve (green) was fit using all data to the form $E_i(N_{ex}) = p(N_{ex})^q + E_i(\infty)$, varying p, q, and $E_i(\infty)$, resulting in an estimate of $E_i(\infty) = 0.36313(7)$ which closely matches the exact $E_{corr} = 0.36320646$ from MOLPRO/MRCC. Fits for fewer data points are denoted $E_i(\infty)$ in the legend, and included data points up to the value of N_{ex} plotted, for varying p, q, and $E_i(\infty)$, as well as varying p and $E_i(\infty)$ and fixing $q = -1$, giving $E_i(\infty) = 0.36317(4)$. The shoulder for this system is at $2.3(3)\times10^5$ excips.

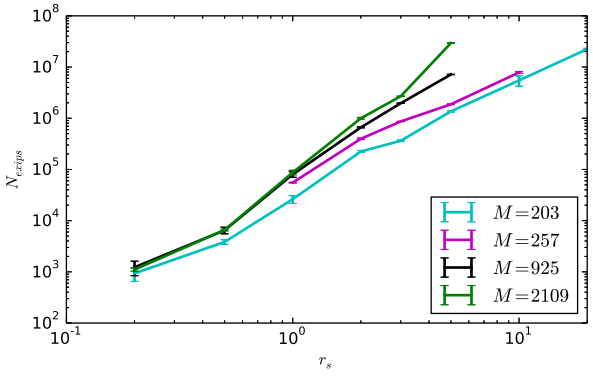

FIG. 8: The shoulder height for CCSDT on the 3D UEG for systems with M plane waves.

than $0.05E_h$. For each population, projected energy estimators for that and lower populations were used in a non-linear least squares fit performed with gnuplot[38]. In Figure 7, we show fits with p and q allowed to vary, as well as fixing $q = -1$. The latter leads to considerably more stable fitting, with data from $N_{ex} < 25000$ providing estimates of $E_i(\infty)$ matching E_{corr} within error bars (albeit large ones), and we have used it as the basis of our later extrapolations. The least squares fits produce error estimates on the values of the fitted variables and we have used these as is as error estimates for the extrapolated values, but caution that they are likely to be under-estimates of the errors.

VII. THE UNIFORM ELECTRON GAS

Whilst the neon atom provides a convenient single-reference system to test the scaling of shoulders with basis, there should be little surprise that the CCSDTQ correlation energy is within a milliHartree of the CCSDT even at the cc-pV6Z basis. In this section, we turn attention to the three dimensional uniform electron gas whose Wigner–Seitz radius, r_s provides a convenient tunable parameter to different regimes of correlation, and which has been little studied with coupled cluster theory. Without extensive parallelism (which is to be the subject of a future paper), we have kept the system sizes small, concentrating on the closed-shell 14-electron gas in plane wave bases. We have used a spherical energy cutoff and denote the number of plane waves M.

r_s	E_{CCSD}	E_{CCSDT}	E_{FCIQMC}
0.1	-0.6639(1)	-0.6659(2)	
0.5	-0.5897(1)	-0.5965(2)	-0.5969(3)
1.0	-0.5155(3)†	-0.5317(3)	-0.5325(4)
2.0	-0.4094(1)	-0.4354(4)‡	-0.4447(4)

TABLE II: Correlation energies for the 14-electron Uniform Electron Gas. Extrapolated initiator stochastic CC results were obtained by basis-set extrapolation from stochastic coupled cluster calculations. FCIQMC results are from Ref 39. † For $r_s = 1$, Shepherd *et al.* have found $E_{\mathrm{CCSD}} = -0.5152(5)$ (Fig 7 of Ref. 41). ‡ Owing to non-converging initiator energies, non-initiator stochastic CC energies were used in the basis set extrapolation.

FCIQMC, CCSD and CCSD(T) results are available in such systems,[39, 40] along with a well-tested methodology to extrapolate to the complete basis set limit[41].

We begin by studying the CCSDT shoulder heights for the UEG as a function of r_s. Figure 8 shows that these rise significantly with r_s, of the order of r_s^3, as well as increasing with the plane wave cutoff energy.

For values of $r_s \leq 2$ initiator CCSD and CCSDT calculations were performed followed by linear fits against $1/M$ for $M \geq 179$, and the $M = \infty$ limit extracted. Results are in Table II. We have also investigated the convergence of the initiator method for these systems, as shown in Figures 9 and 10. Whilst the convergence of the extrapolation for $r_s = 1$ in Figure 9 appears to well-match the exact result, we find the same methodology predicts a result many standard deviations away from the true correlation energy for $r_s = 2$ (Figure 10). This non-convergent behaviour of the initiator approximation may be owing to changes in the nature of the correlation with increasing r_s, as the UEG is known to become more multireference with increasing r_s. We therefore would caution against the use of initiator extrapolation in such cases where there is not a clear sign of convergence.

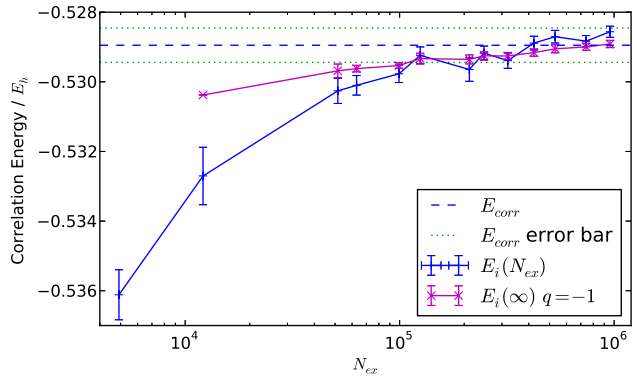

FIG. 9: CCSDT initiator extrapolation for the 14-electron 3D uniform electron gas with $M = 925$ and $r_s = 1$.

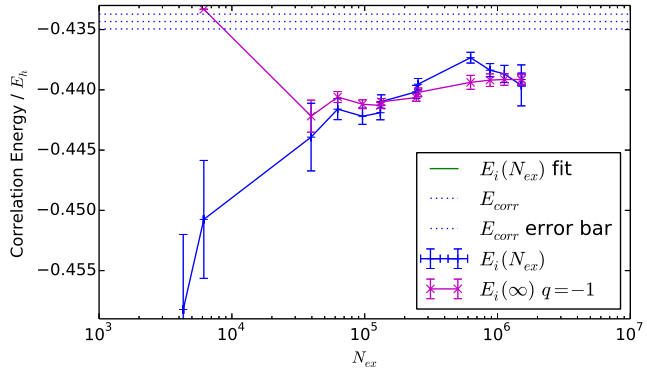

FIG. 10: CCSDT initiator extrapolation for the 14-electron 3D uniform electron gas with $M = 925$ and $r_s = 2$. Though the raw initiator energies do not appear to be converging, the extrapolation does appear to converge to a lower energy than the true correlation energy, E_{corr}, which was calculated with non-initiator CCSDTMC.

VIII. CONCLUSION

We have investigated the stability of Stochastic Coupled Cluster calculations and suggest that the particle ratio (between the total excip population and the population on the reference) is a convenient measure of the progress of a Stochastic Coupled Cluster calculation. The particle ratio exhibits a maximum as a calculation progresses through the growth phase, and this maximum signals the critical number of particles required for a stable calculation, and indicates a 'shoulder height' equivalent to the plateau height in FCIQMC. As a fraction of the truncated Hilbert space of a coupled cluster calculation, this effective plateau height remains a similar order of magnitude (1–10%) as basis set is increased, and decreases as the truncation level increases, so allowing an estimate of the upper bound storage required of a CCMC calculation.

We have investigated two possible implementations of the initiator approximation in CCMC, and of the most stable, we conclude that the initiator approximation performs a similar function to that in FCIQMC, and stabilizes calculations with low numbers of excips at the expense of introducing a systematic initiator error. We have demonstrated that this error, which in all cases we have investigated causes a systematic lowering of the energy, decays as N_{excip}^{-1} and this can be used as a basis for extrapolation of the true correlation energy. Unlike FCIQMC, we have found that the initiator error does not decay away quickly in comparison to the critical stable population indicated by 'shoulder plots'. In both the 3D Uniform Electron Gas and Ne atoms in large bases, we show that the particle population require to reduce the initiator error sufficiently is of the order of the shoulder height. Therefore the reductions in particle population, which in FCIQMC can amount to many orders of magnitude, do not appear to be present in Stochastic Coupled Cluster Theory.

However, the benefits of the initiator approximation are still manifold. With careful extrapolation, it is possible to perform calculations which are not possible without use of the initiator approximation. Additionally, the initiator approximation dramatically reduces the variance of the projected energy within a calculation, and thus by its use the required calculation times are greatly reduced. There is reason to be optimistic as there is much which can be tuned in the Stochastic Coupled Cluster algorithm. In particular, the role of the population of particles on the reference, which acts as a normalization, is critical in the stability of calculations, and in the projected energy. With further algorithmic modifications and optimization of parameters, the convergence of the initiator energy with particle number will likely become faster allowing yet larger calculations to be performed.

Finally we demonstrate that this method is capable of calculating coupled cluster energies of the Uniform Electron Gas up to $r_s = 2$. Studies beyond this will require significantly more computational resources, but are still possible. Of particular note is the vital contribution of triple excitations even at these modest values of r_s, and we hope to make a more detailed study of this in the future.

Acknowledgments

We are grateful to Miss Ruth Franklin and Mr William Vigor for discussions on the algorithms. AJWT thanks the Royal Society for a University Research Fellowship and Imperial College London for a Junior Research Fellowship, where this work was started. JSS acknowledges the research environment provided by the Thomas Young Centre under Grant No. TYC-101. Calculations were performed on the Imperial College High Performance Computing facilities[42] and the University of Cambridge High Performance Computing facilities using

the NECI[43] and the HANDE[44, 45] codes. Molecular Orbital integrals were generated with a modified version of Q-Chem[46], and for Ne cc-pV6Z, MOLPRO[35]. Raw and analysed data is available in Ref. 47. We use an automated iterative blocking algorithm[34, 48–50] to accurately estimate the stochastic error in all CCMC calculations presented in this paper. Figures were plotted using matplotlib[51].

Appendix A: Initiator Extrapolation

We derive a form for the convergence of the initiator approximation with psip population in Full Configuration Interaction Quantum Monte Carlo. We expect that, while CCMC behaviour will be more complicated, it should follow the same general form.

Let us denote the true ground state wavefunction ψ_0 with energy E_0 and impose intermediate normalization on it such that $\langle D_{\mathbf{0}}|\psi_0\rangle = 1$. Consider the instantaneous wavefunction at a given timestep, $\psi(\tau)$. We can expect it to consist of a some amount of the ground state, and a difference term, $\psi(\tau) = C(\tau)\psi_0 + \Delta(\tau)$, where we can specify that $\langle D_{\mathbf{0}}|\Delta(\tau)\rangle = 0$. In a converged simulation C is the population of the reference determinant, and is proportional to the total number of particles. The projected energy is given by

$$E(\tau) = \frac{\langle D_{\mathbf{0}}|\hat{H}|\psi(\tau)\rangle}{\langle D_{\mathbf{0}}|\psi(\tau)\rangle} = \frac{\langle D_{\mathbf{0}}|\hat{H}|\psi(\tau)\rangle}{C(\tau)}. \tag{A1}$$

In an unbiased FCIQMC calculation $\langle\Delta(\tau)\rangle_\tau = 0$ (where $\langle\rangle_\tau$ indicates the long time average over τ). In moving to the initiator modification, we introduce a systematic bias in the dynamics, resulting in a $\langle\Delta(\tau)\rangle_\tau$ which no longer vanishes. For a given value of the initiator threshold and value of C, there will be a maximum component of $\Delta(\tau)$, i.e. some δ, such that for all determinants $\mathbf{i}$, $\langle D_{\mathbf{i}}|\Delta(\tau)\rangle < \delta$. Let us assume that any dependence δ has on C is at most of order C^0 (i.e. constant), as a higher power would lead to the initiator approximation not converging with increasing psip population.

Explicitly evaluating the projected energy,

$$\begin{aligned} E(\tau) &= \frac{\langle D_{\mathbf{0}}|\hat{H}|C(\tau)\psi_0 + \Delta(\tau)\rangle}{C(\tau)} && \text{(A2)}\\ &= C(\tau)\frac{\langle D_{\mathbf{0}}|\hat{H}|\psi_0\rangle}{C(\tau)} + \frac{\sum_{\mathbf{i}\neq\mathbf{0}}\langle D_{\mathbf{0}}|\hat{H}|D_{\mathbf{i}}\rangle\langle D_{\mathbf{i}}|\Delta(\tau)\rangle}{C(\tau)} && \text{(A3)}\\ &\leq E_0 + \frac{\sum_{\mathbf{i}} H_{0\mathbf{i}}\delta}{C(\tau)}. && \text{(A4)} \end{aligned}$$

Denoting the maximal value of $H_{0\mathbf{i}}$ as η, if there are N_{doub} double excitations of the reference, we find

$$|E(\tau) - E_0| \leq \frac{N_{\text{doub}}\eta\delta}{C(\tau)}. \tag{A5}$$

Since N_{doub} is constant, and $C(\tau)$ scales linearly with N_{psip}, we may attempt to fit to $(N_{\text{psip}})^{-1}$ as an upper bound.

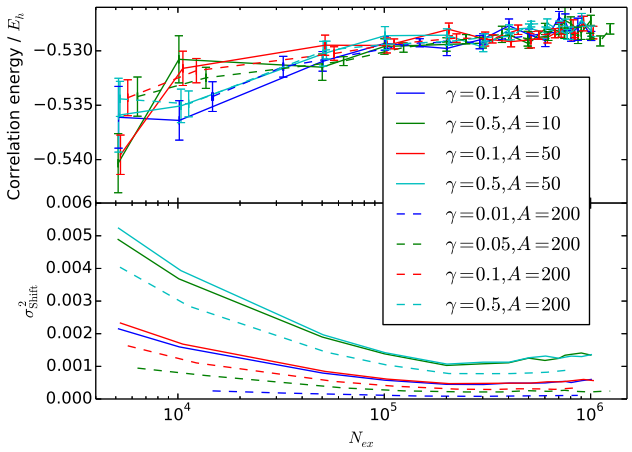

FIG. 11: iCCMC convergence studies of CCSDT for the 14-electron 3D uniform electron gas at $r_s = 1$ using $M = 925$ plane waves for different population control parameters (top) and the corresponding variance in the shift observed during each calculation (bottom). $\delta\tau = 10^{-4}$ in all calculations.

Appendix B: Population Control Effects

Appendix A contains the assumption that the initiator approximation is the cause of all systematic errors in the simulation. This is not strictly true; the effect of population control is well understood to introduce a small bias which decays with increasing population[30, 52]. We briefly investigate here the effect of population control on the convergence of the initiator error.

The shift is updated every A timesteps during the population control phase according to[9, 30]

$$S(\tau) = S(\tau - A\delta\tau) - \frac{\gamma}{A\delta\tau}\ln\left(\frac{N_{ex}(\tau)}{N_{ex}(\tau - A\delta\tau)}\right), \tag{B1}$$

where γ is a (typically small) damping parameter. The variance in the shift (which is independent of the length of a simulation) is a measure of the population fluctuations caused by requiring the population to be approximately constant. Figs. 11 and 12 show the effect of population control on the 14-electron UEG at $r_s = 1$. All simulations converge to within stochastic error at sufficiently large populations (Fig. 11); however extrapolation also depends upon data from simulations using smaller populations of excips. Whilst this is far from a comprehensive study, we observe a relatively smooth convergence in all cases once $\sigma^2_{\text{Shift}} < 0.002$ (Fig. 12). Therefore extrapolation requires careful choices of these parameters.

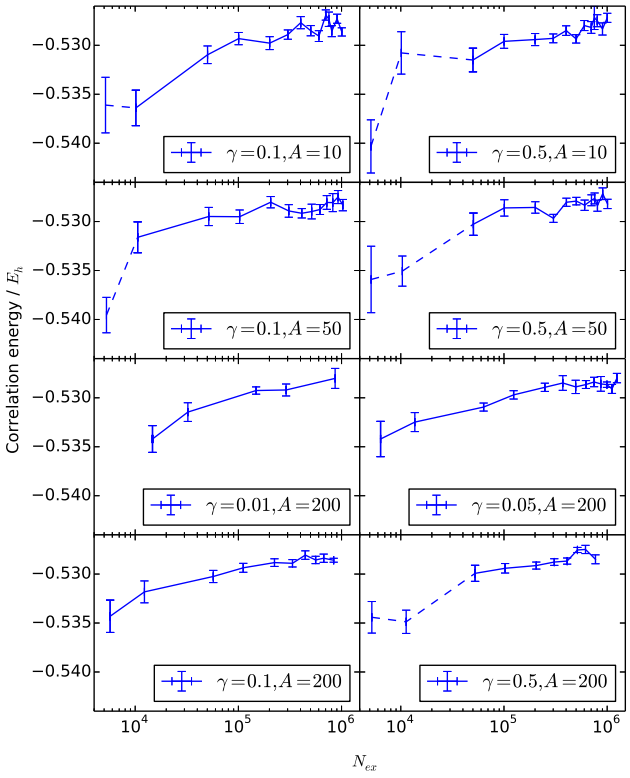

FIG. 12: iCCMC convergence studies of CCSDT for the 14-electron 3D uniform electron gas at $r_s = 1$ using $M = 925$ plane waves for different population control parameters. $\delta\tau = 10^{-4}$ in all calculations. Dashed lines represent calculations for which the shift variance is greater than 0.002. This is a more detailed view of the top panel in Fig. 11.

[1] A. J. W. Thom, Phys. Rev. Lett. **105**, 263004 (2010).
[2] J. Čížek, J. Chem. Phys. **45**, 4265 (1966).
[3] J. Čížek and J. Paldus, Int. J. Quantum Chem. **5**, 359 (1971).
[4] R. J. Bartlett and M. Musiał, Rev. Mod. Phys. **79**, 291 (2007).
[5] K. Raghavachari, G. W. Trucks, J. A. Pople, and M. Head-Gordon, Chem. Phys. Lett. **157**, 479 (1989).
[6] D. G. Liakos, M. Sparta, M. K. Kesharwani, J. M. L. Martin, and F. Neese, J. Chem. Theory Comput. **11**, 1525 (2015).
[7] J. E. Subotnik, A. Sodt, and M. Head-Gordon, J. Chem. Phys. **125**, 074116 (2006).
[8] M. Schütz, Phys. Chem. Chem. Phys. **4**, 3941 (2002).
[9] G. H. Booth, A. J. W. Thom, and A. Alavi, J. Chem. Phys. **131**, 054106 (2009).
[10] G. H. Booth and A. Alavi, J. Chem. Phys. **132**, 174104 (2010).
[11] D. Cleland, G. H. Booth, and A. Alavi, J. Chem. Phys. **132**, 41103 (2010).
[12] G. H. Booth, D. Cleland, A. J. W. Thom, and A. Alavi, J. Chem. Phys. **135**, 084104 (2011).
[13] D. Cleland, G. H. Booth, C. Overy, and A. Alavi, J. Chem. Theor. Comput. **8**, 4138 (2012).
[14] G. H. Booth, D. Cleland, A. Alavi, and D. P. Tew, J. Chem. Phys. **137**, 1 (2012).
[15] C. Overy, G. H. Booth, N. S. Blunt, J. J. Shepherd, D. Cleland, and A. Alavi, J. Chem. Phys. **141**, 1 (2014).
[16] R. E. Thomas, D. Opalka, C. Overy, P. J. Knowles, A. Alavi, and G. H. Booth, J. Chem. Phys. **143**, 1 (2015).
[17] N. S. Blunt, A. Alavi, and G. H. Booth, Phys. Rev. Lett. **115**, 050603 (2015).
[18] N. S. Blunt, S. D. Smart, G. H. Booth, and A. Alavi, J. Chem. Phys. **143**, 134117 (2015).
[19] J. J. Shepherd, A. Grüneis, G. H. Booth, G. Kresse, and A. Alavi, Phys. Rev. B **86**, 035111 (2012).
[20] J. J. Shepherd, G. Booth, A. Grüneis, and A. Alavi, Phys. Rev. B **85**, 081103 (2012).
[21] J. J. Shepherd, G. E. Scuseria, and J. S. Spencer, Phys. Rev. B **90**, 155130 (2014).
[22] A. Grüneis, G. H. Booth, M. Marsman, J. Spencer, A. Alavi, and G. Kresse, J. Chem. Theor. Comput. **7**, 2780 (2011).
[23] G. H. Booth, A. Grüneis, G. Kresse, and A. Alavi, Nature

493, 365 (2013).
[24] F. R. Manby, D. Alfè, and M. J. Gillan, Phys. Chem. Chem. Phys. **8**, 5178 (2006).
[25] M. Marsman, A. Grüneis, J. Paier, and G. Kresse, J. Chem. Phys. **130**, 184103 (2009).
[26] A. Grüneis, J. Chem. Phys. **143**, 102817 (2015).
[27] M. Del Ben, J. Hutter, and J. VandeVondele, J. Chem. Phys. **143**, 102803 (2015).
[28] A. Grüneis, Phys. Rev. Lett. **115**, 066402 (2015).
[29] J. B. Anderson, J. Chem. Phys. **63**, 1499 (1975).
[30] C. J. Umrigar, M. P. Nightingale, and K. J. Runge, J. Chem. Phys. **99**, 2865 (1993).
[31] J. S. Spencer, N. S. Blunt, and W. M. Foulkes, J. Chem. Phys. **136**, 054110 (2012).
[32] T. H. Dunning, Jr., J. Chem. Phys. **90**, 1007 (1989).
[33] A. K. Wilson, T. van Mourik, and T. H. Dunning, Jr., J. Mol. Struct. (THEOCHEM) **388**, 339 (1997).
[34] H. Flyvbjerg and H. G. Petersen, J. Chem. Phys. **91**, 461 (1989).
[35] H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, M. Schütz, P. Celani, W. Györffy, D. Kats, T. Korona, R. Lindh, et al., *Molpro, version 2010.2, a package of ab initio programs* (2010), see http://www.molpro.net.
[36] H.-J. Werner, P. J. Knowles, G. Knizia, F. R. Manby, and M. Schütz, WIREs Comput Mol Sci **2**, 242 (2012).
[37] M. Kállay, Z. Roli, J. Csonto, I. Ladjánszk, L. Szeged, B. Ladóczk, and G. Samu, *MRCC, a quantum chemical program suite*, `http://www.mrcc.hu`.
[38] T. Williams, C. Kelley, and many others, *Gnuplot 4.6: an interactive plotting program*, `http://www.gnuplot.info` (2013).
[39] J. J. Shepherd, G. H. Booth, and A. Alavi, J. Chem. Phys. **136**, 244101 (2012).
[40] J. J. Shepherd, T. M. Henderson, and G. E. Scuseria, Phys. Rev. Lett. **112**, 133002 (2014).
[41] J. J. Shepherd, A. Grüneis, G. Booth, G. Kresse, and A. Alavi, Phys. Rev. B **86**, 035111 (2012).
[42] `http://www.imperial.ac.uk/ict/services/teachingandresearchservices/highperformancecomputing`.
[43] G. H. Booth, S. D. Smart, and A. Alavi, Mol. Phys. **112**, 1855 (2014).
[44] `http://www.hande.org.uk`.
[45] J. S. Spencer, N. S. Blunt, W. A. Vigor, F. D. Malone, W. M. C. Foulkes, J. J. Shepherd, and A. J. W. Thom, J. Open Research Software **3** (2015), `http://doi.org/10.5334/jors.bw`.
[46] Y. Shao, Z. Gan, E. Epifanovsky, A. T. Gilbert, M. Wormit, J. Kussmann, A. W. Lange, A. Behn, J. Deng, X. Feng, et al., Mol. Phys. **113**, 184 (2015).
[47] J. Spencer and A. J. W. Thom, *Data from "developments in stochastic coupled cluster theory: The initiator approximation and application to the uniform electron gas"* (2015), URL `http://dx.doi.org/10.5281/zenodo.33883`.
[48] U. Wolff, Comput. Phys. Commun. **156**, 143 (2004).
[49] R. M. Lee, G. J. Conduit, N. Nemec, P. López Rios, and N. D. Drummond, Phys. Rev. E **83**, 066706 (2011).
[50] See `https://github.com/jsspencer/pyblock` for code.
[51] J. D. Hunter, Computing In Science & Engineering **9**, 90 (2007).
[52] W. A. Vigor, J. S. Spencer, M. J. Bearpark, and A. J. W. Thom, J. Chem. Phys. **142**, 104101 (2015).
[53] More formally $\mathcal{O}(N^3M^4)$ where M is the number of basis functions in the system
[54] We have adopted $\delta\tau$ rather than τ as in [1] so as to be more consistent with the notation of Alavi *et al.*
[55] The converse, that solutions to the iteration will also be Coupled Cluster solutions, is not necessarily true, though we have not found any situations where this is the case. Indeed there is no guarantee that this iteration procedure will converge, though with a sensible choice of $\delta\tau$ it appears to be convergent for primarily single-reference systems
[56] Note that our previous publication used the word excitor for both the operator and the particle.
[57] This makes use of the PELE suite (see `https://github.com/pele-python/pele`), and we will describe details of this elsewhere